Collaborative High Performance Computing Course Using the Access Grid

Roger Edberg, Guy Robinson
Arctic Region Supercomputing Center, University of Alaska Fairbanks

Jennifer Parham, Don Morton
Montana Rockies Center for Computational Science, University of Montana

Jeff Shuckra, Tim Warburton, Frank Gilfeather
High Performance Computing, Education and Research Center, University of New Mexico

Access Grid Retreat 2002, San Diego CA, 4-5 March 2002

Overview

In 2001 the Arctic Region Supercomputing Center and Physics Department at the University of Alaska Fairbanks partnered with the Computer Science departments at the University of Montana and the University of New Mexico to offer a collaborative, for-credit, graduate-level course in parallel computing via the Access Grid. The three institutions used the Access Grid to supplement a standard schedule of course lectures, tutorials and lab work, including at least one weekly AG meeting of all students and instructors.

The idea to work together via the AG to teach a high-performance computing course came out of discussions between Alaska, Montana and New Mexico at the Access Grid Tutorial 2000 and the AG Retreat 2001. Our Access Grid Nodes could be used to supplement courses in high-performance computing at the three universities, share our teaching expertise and potentially provide a richer environment for students and instructors. The class, scheduled for Fall Semester 2001, was listed differently in the course catalogs at the three institutions to meet local administrative and curriculum requirements:

Site	Course Number/Title	Instructors	Instructor Expertise
UAF	PHYS693: Parallel Scientific Computation	Guy Robinson	Computational Physics, MPP systems
UMT	CS495: High Performance Computing for Scientists	Jennifer Parham, Don Morton	Computer Science, MPP applications
UNM	CS471/ME471: Introduction to Parallel Programming	Tim Warburton, Brian Smith	Computer Science

The common course thread was (applied) parallel computing; the target audience was science or engineering students at the upper-undergraduate or graduate level. The primary instructors at each institution actually had slightly different educational and research backgrounds, some with a computer science focus and others with computational science experience.

The designated AG Class meeting time was on Thursdays, 12-2pm AK time or 2-4pm Mountain Time. Instructors and operations staff held a weekly teleconference on Wednesdays to confirm and review the AG session agenda, prepare materials, and discuss the progress of the class. Presentation and lecture duties rotated between the instructors.

ARSC, Montana and New Mexico also worked together to produce and present a program about the course for SC Global, at SC2001. In addition to weekly course meetings all three sites took part in an intensive schedule of test cruises and preparation work for SC Global from August through November. During most of September, October, and early-November 2001 ARSC was online 8-12 hours per week on the Access Grid.

Course Experiences

Our experiences preparing and delivering the collaborative class are summarized by topic below.

1. Students and instructors did benefit from exposure to multiple points of view, differing areas of expertise, and exposure to different HPC resources via course lectures and demonstrations. In this regard AG technology did complement the traditional class/laboratory teaching setting.

2. Lectures and demonstrations often went in unexpected directions, enabled by the Access Grid interaction and access to additional resources, for example, HPC and programming documentation on the web. AG technology enabled spontaneous discussion and research during class. There were times when the class developed the atmosphere of a field trip.

3. The class experience proved to be new, interesting and sometimes fun for all involved. A good-natured rapport developed between the student groups and instructors. Round-robin AG demonstrations involving the students and instructors were used to illustrate data communications and parallel processing concepts.

4. Access Grid technology is still cumbersome to work with and requires several staff to maintain and use properly. The Alaska, Montana and New Mexico AG teams spent a LOT of time dealing with technical issues, which detracted from being able to teach the course. Audio problems were particularly disruptive; telco backup was next to useless on at least one occasion. Technical difficulties did become less frequent and serious during the semester as we gained experience, and were taken in stride after the first classes. We never had to completely fall back to "offline" mode for any classes.

5. Instructors faced a new set of challenges. They had to become familiar with Access Grid presentation style and protocol, and the new and unfamiliar "AG classroom" environment. All of the instructors had to work out new strategies for engaging remote students, gaining their trust, encouraging discussion, and tightening the teacher-student feedback loop in the new environment. Instructors also had to work harder to stay synchronized with the course plan. One instructor commented that he felt intensely frustrated that he couldn't just grab some chalk, go to a blackboard and start writing and sketching to illustrate ideas. The workload for instructors and operations staff was much heavier than anticipated.

6. The single "virtual classroom" did not materialize. End-of-course evaluations indicated two reasons for this: (1) class projects were assigned to be local rather than cross-institutional via the Access Grid, and (2) the AG required node operations staff, who were not available at all hours, and at all sites. AG connectivity was generally very good, however it wasn't seamless enough to produce the illusion that everybody was in the same classroom. Audio and network problems sometimes detracted from the classroom experience. There were also instances when one or more sites could not interact because of network multicast problems. On the other hand the three groups did begin to feel comfortable with one another, especially toward the end of the semester after earlier audio and network problems were resolved. Students became more at ease and felt more comfortable participating as they gained experience. Interaction between only two sites might better approach the ideal of a "virtual classroom." Groups of students might interact better in a longer series of AG classes.

7. Access Grid technology did not magically bridge distance. The AG environment requires additional components to support student-student interactions. Currently, student conversation and comment has to be facilitated and actively encouraged; students must wait their turn to speak among themselves or to other peers at remote nodes. A student-only MUD might be one way to increase "live" interaction without interrupting the main lecture or demonstration. A mail reflector was used to keep the students informed of class schedules, resource materials and for questions to the instructors.

8. Social and environmental factors are important. Room layout, camera placement, and presentation style affect the quality of classroom experience for everyone. Instructors learned with experience to speak clearly and look into the camera rather than speaking to the AG display screen. Instructors benefited from practice time or at least some "live" AG experience before being turned loose in the classroom.

9. Planning and coordination on course content, presentation and AG operations are critical. The weekly AK-MT-NM class session typically involved at least three node operators, two or three instructors, approximately 10-15 students and additional operations staff as required. Simply keeping everybody informed was a major chore. Email worked well for this but could be unreliable if any one person did not continuously check their mail. Telephone was much more effective and reliable for important notices. (Recommendation: use the phone or teleconferences for important discussions; involve everybody.) The level of orchestration required was a lot higher than anticipated.

10. Multi-institution Access Grid teaching projects need to fit the administrative framework of each partner. Mismatches in starting dates and length of academic semesters or quarters can be problematic. Fall academic term began on different dates for our three universities, with as much as two weeks of lag time before all three classes could meet as a unified group. There were similar differences in semester breaks and final exam timetables. Curriculum requirements have to overlap. Students were accountable only to the local course instructor. Instructors did not evaluate or give exams to remote students.

11. Time zone differences are important. Class times should be planned such that no group should have to work at inopportune or inconvenient times, unless they choose to. The Alaska to Mountain time zone difference was two hours, which presented no difficulties for any of the sites. ARSC also worked with AG sites in Asia and the Pacific Rim (primarily in Australia, China and Japan) test cruising for SC Global. Morning in Beijing, Sydney or Tokyo corresponded to late afternoon Alaska time, and late evening US Central or Eastern time.

12. Dedicated, Access Grid-enabled classroom space is essential for a collaborative course. ARSC's AG node space is shared with other departments on the UAF campus. There were (and will continue to be) times when ARSC was unable to participate in AG class sessions because we did not have access to our facilities.

13. Tools like VNC and the Mimeo Electronic Whiteboard added a new dimension to class demonstrations and lectures. Students and instructors were, generally speaking, happy with these tools after they become comfortable with them. VNC demonstrations using VAMPIR to profile parallel codes were successful and well accepted by students. VNC let the instructors use familiar HPC tools and present them to the class, without a lot of additional learning effort or operational overhead.

14. Murphy's Law scales non-linearly with the number of Access Grid nodes and the number of people involved. Less than a week before SC2001, a backhoe cut through a critical network trunk line cable in Anchorage, leaving ARSC unable to test AG connectivity until Monday of SC2001 week. Staff schedules and commitments changed at short notice. Endless PowerPoint revisions circulated among our group while nobody knew which was the final version. Email from ARSC staff in Denver had to go back via Fairbanks to reach other people in Denver until local network problems were fixed. Luck was also on our side. SciNet suddenly stabilized an hour or so before our presentation, after a day of continual problems. ARSC made a timely comeback into the main SC Global presentation venue, just as it's speaker was scheduled to come on. The SC Global presentation came off without any major problems. The bottom line: be prepared for anything! A positive, flexible outlook is essential.

15. The NCSA Scheduler is a good tool that needs some fine tuning to make it a great centralized resource for the Access Grid community. ARSC staff had problems using the Scheduler via Macintosh machines before and during SC2001. The Scheduler and other AG components should be completely platform-neutral. Mac users are currently second class citizens with respect to AG-ware.

16. Some SC Global requirements were orthogonal to the class requirements and plan. SC Global was scheduled on days where the group class did not meet, so we were unable to show the class in action. SC2001 happened halfway through Fall Semester. Our SC Global presentation had to be at least in part a mid-term review of the class and how the AG collaboration had worked. A more complete review and evaluation would happen at the conclusion of the semester. SC Global did create some extracurricular work and distraction for students and the instructors. On the other hand, the continual series of SC Global test cruises helped us fix technical problems, gave us valuable practice, and improved our AG competence. Time zone differences (item 11 above) were a factor here as well. ARSC maintained a steady schedule of 6:00 am(!) and 4:00 pm AK time test cruises through September and October 2001.

17. Access Grid "trial by fire" can be good, provided stress and expectations are kept reasonable. The class and SC Global forced us all to work together, keep our AG nodes up and running, and do the work and gain the experience we needed. We at ARSC felt like we progressed from being Access Grid beginners to functional, competent members of the community.

Many of the issues described above scale with the complexity of the collaboration and the number of sites involved. A two-way AG teaching collaboration would probably run more smoothly than our three-way model. Each participant site adds another layer of content, presentation, node operations and management priorities to the mix.

The success of this type of teaching collaboration depends on the priorities and focus of the participating groups. The project framework has to suit all of the parties. The required effort will probably frustrate groups with only a marginal interest, or those looking to do this sort of project part-time. Sites must also be able to dedicate full-time staff to the project.

Conclusion and Future Directions

ARSC, Montana and New Mexico plan to continue using the AG as a supplement to teaching, as a vehicle for bringing guest speakers and course-relevant demonstrations to the student audience. Members of ARSC staff are lined up as guest lecturers for the continuation of the high-performance computing course during Spring Semester 2002.

ARSC and Montana have discussed the idea of running an intensive AG workshop on Scientific Visualization, possibly several days to a week in length. Discussions are continuing on dates, content and possible formats for this workshop. ARSC is also using its AG Node for collaborative visualization, virtual environments and computer art, as well as for administrative meetings with other HPC centers and agencies.

Montana is, in collaboration with Motorola Labs, planning to pursue initiatives that lead to better understanding of how we can make the Access Grid truly accessible to the general population in the context of supporting Real Life activities, such as project-based collaborations.

The University of New Mexico is promoting the installation of 20 AG nodes throughout the state to educational facilities and indigenous tribes. They are currently conducting workshops and technology assessments to teach the steps of installing network and programming. The Tribal Virtual Network (TVN) is implementing five AG nodes and is looking forward to working with UNM's Extended University to deliver college accredited courses.

The Access Grid proved to be a valuable supplement to the traditional classroom teaching environment. On the other hand, all of those involved in our project also feel that we did not really come close to fully exploiting the potential of the AG for teaching and research collaboration. AG technologies are certain to become more widely accessible, more portable, more user friendly, and less expensive in the future. More important, they are likely to create

new means of interaction. There is huge potential to explore and refine teaching collaboration strategies.

Future directions for using the Access Grid in a teaching environment must include tighter integration of AG tools with the native fabric of the course, and whatever resources are required for the subject. In our case this would ideally be seamless integration with HPC resources and tools: MPP systems, HPC clusters, large-scale data storage and archives, and visualization and VE systems. Students of high performance computing with access to these resources should have a richer learning experience.

The Team

Node Operators:
Bob Huebert, Steve Munk, Paul Mercer (Alaska)
Jeremy Sauer, Don Morton (Montana)
Jeff Shuckra (New Mexico)

Course Instructors:
Guy Robinson (Alaska)
Jennifer Parham, Don Morton (Montana)
Tim Warburton, Brian Smith (New Mexico)

SC Global Coordination:
Roger Edberg, Jim Long (Alaska)

Management and Vision:
Franks Williams, Director, ARSC/Univ. of Alaska Fairbanks (UAF)
Barbara Horner-Miller, Director-User Services, ARSC/UAF
Virginia Bedford, Director-Technical Services, ARSC/UAF
Ray Ford, Associate Vice President for Information Technology, Univ. of Montana
Thomas Storch, Dean-College of Arts and Sciences, Univ. of Montana
Jerry Esmay, Chair-Department of Computer Science, Univ. of Montana
Frank Gilfeather, Director-HPCERC, Univ. of New Mexico

The authors would like to thank Terry Disz, Mary Fritsch, Bob Olson, Julie Wulf, Jennifer Teig von Hoffman and, Karen Camarda and her University of Kansas Nodeop team for their help in making our SC Global presentation a success.